

ATTORNEY DOCKET No.

NVIDP008/P000168

U.S. PATENT APPLICATION
FOR
A TRANSFORM, LIGHTING AND
RASTERIZATION SYSTEM EMBODIED ON A
SINGLE SEMICONDUCTOR PLATFORM

Inventors

Erik Lindholm
Simon Moy
Kevin Dawallu
Mingjian Yang
John Montrym
David B. Kirk
Paolo Sabella
Matthew N. Papakipos
Douglas A. Voorhies
Nicholas J. Foskett

ASSIGNEE: nVIDIA CORPORATION

SILICON VALLEY IP GROUP
P.O. BOX 721120
SAN JOSE, CA 95172-1120
TELEPHONE: (408) 971-2573

A TRANSFORM, LIGHTING AND RASTERIZATION SYSTEM EMBODIED ON A SINGLE SEMICONDUCTOR PLATFORM

5

RELATED APPLICATIONS

The present application is related to applications entitled “Method, Apparatus and Article of Manufacture for Area Rasterization using Sense Points” which was filed under attorney docket number NVIDP005, “Method, Apparatus and Article of

10 Manufacture for Boustrophedonic Rasterization” which was filed under attorney docket number NVIDP006, “Method, Apparatus and Article of Manufacture for Clip-less Rasterization using Line Equation-based Traversal” which was filed under attorney docket number NVIDP007, “Method, Apparatus and Article of Manufacture for a Vertex Attribute Buffer in a Graphics Processor” which was filed under attorney docket number NVIDP009, “Method, Apparatus and Article of Manufacture for a Transform Module in a Graphics Processor” which was filed under attorney docket number NVIDP010, ”Method and Apparatus for a Lighting Module in a Graphics Processor” which was filed under attorney docket number NVIDP011, and “Method, Apparatus and Article of Manufacture for a Sequencer in a

15

20 Transform/Lighting Module Capable of Processing Multiple Independent Execution Threads” which was filed under attorney docket number NVIDP012 which were filed concurrently herewith, and which are all incorporated herein by reference in their entirety.

25

FIELD OF THE INVENTION

The present invention relates generally to graphics processors and, more particularly, to graphics pipeline systems including transform, lighting and rasterization modules.

30

BACKGROUND OF THE INVENTION

Three dimensional graphics are central to many applications. For example, computer aided design (CAD) has spurred growth in many industries where 5 computer terminals, cursors, CRT's and graphics terminals are replacing pencil and paper, and computer disks and tapes are replacing drawing vaults. Most, if not all, of these industries have a great need to manipulate and display three-dimensional objects. This has lead to widespread interest and research into methods of modeling, rendering, and displaying three-dimensional objects on a computer screen or other 10 display device. The amount of computations needed to realistically render and display a three-dimensional graphical object, however, remains quite large and true realistic display of three-dimensional objects have largely been limited to high end systems. There is, however, an ever-increasing need for inexpensive systems that can quickly and realistically render and display three dimensional objects.

15 One industry that has seen a tremendous amount of growth in the last few years is the computer game industry. The current generation of computer games is moving to three-dimensional graphics in an ever increasing fashion. At the same time, the speed of play is being driven faster and faster. This combination has fueled 20 a genuine need for the rapid rendering of three-dimensional graphics in relatively inexpensive systems. In addition to gaming, this need is also fueled by e-Commerce applications, which demand increased multimedia capabilities.

Rendering and displaying three-dimensional graphics typically involves 25 many calculations and computations. For example, to render a three dimensional object, a set of coordinate points or vertices that define the object to be rendered must be formed. Vertices can be joined to form polygons that define the surface of the object to be rendered and displayed. Once the vertices that define an object are 30 formed, the vertices must be transformed from an object or model frame of reference to a world frame of reference and finally to two-dimensional coordinates that can be displayed on a flat display device. Along the way, vertices may be rotated, scaled,

eliminated or clipped because they fall outside the viewable area, lit by various lighting schemes, colorized, and so forth. Thus the process of rendering and displaying a three-dimensional object can be computationally intensive and may involve a large number of vertices.

5

A general system that implements such a pipelined system is illustrated in Prior Art Figure 1. In this system, data source 10 generates a stream of expanded vertices defining primitives. These vertices are passed one at a time, through pipelined graphic system 12 via vertex memory 13 for storage purposes. Once the 10 expanded vertices are received from the vertex memory 13 into the pipelined graphic system 12, the vertices are transformed and lit by a transformation module 14 and a lighting module 16, respectively, and further clipped and set-up for rendering by a rasterizer 18, thus generating rendered primitives that are displayed on display device 20.

15

During operation, the transform module 14 may be used to perform scaling, rotation, and projection of a set of three dimensional vertices from their local or model coordinates to the two dimensional window that will be used to display the rendered object. The lighting module 16 sets the color and appearance of a vertex 20 based on various lighting schemes, light locations, ambient light levels, materials, and so forth. The rasterization module 18 rasterizes or renders vertices that have previously been transformed and/or lit. The rasterization module 18 renders the object to a rendering target which can be a display device or intermediate hardware or software structure that in turn moves the rendered data to a display device.

25

When manufacturing graphics processing systems, there is a general need to increase the speed of the various graphics processing components, while minimizing costs. In general, integration is often employed to increase the speed of a system. Integration refers to the incorporation of different processing modules on a single 30 integrated circuit. With such processing modules communicating in a microscopic semiconductor environment, as opposed to external buses, speed is vastly increased.

Integration if often limited, however, by a cost of implementing and manufacturing multiple processing modules on a single chip. In the realm of graphics processing, any attempt to integrate the transform, lighting, and

5 rasterization modules for increased speed would be cost prohibitive. The reason for this increase in cost is that the required integrated circuit would be of a size that is simply too expensive to be feasible.

This size increase is due mainly to the complexity of the various engines.

10 High performance transform and lighting engines alone are very intricate and are thus expensive to implement on-chip, let alone implement with any additional functionality. Further, conventional rasterizers are multifaceted with the tasks of clipping, rendering, etc. making any cost-effective attempt to combine such module with the transform and lighting modules nearly impossible.

15 There is therefore a need for a transform, lighting, and rasterization module having a design that allows cost-effective integration.

DISCLOSURE OF THE INVENTION

A graphics pipeline system is provided for graphics processing. Such system includes a transform module adapted for being coupled to a vertex attribute buffer for receiving vertex data. The transform module serves to transform the vertex data from object space to screen space. Coupled to the transform module is a lighting module which is positioned on the single semiconductor platform for performing lighting operations on the vertex data received from the transform module. Also included is a rasterizer coupled to the lighting module and positioned on the single semiconductor platform for rendering the vertex data received from the lighting module.

In one aspect of the present invention, the transform module is designed to facilitate integration by including an input buffer adapted for being coupled to a vertex attribute buffer for receiving vertex data therefrom. A multiplication logic unit has a first input coupled to an output of the input buffer. Also provided is an arithmetic logic unit having a first input coupled to an output of the multiplication logic unit. Coupled to an output of the arithmetic logic unit is an input of a register unit.

An inverse logic unit is also provided including an input coupled to the output of the arithmetic logic unit for performing an inverse or an inverse square root operation. Further included is a conversion module coupled between an output of the inverse logic unit and a second input of the multiplication logic unit. In use, the conversion module serves to convert scalar vertex data to vector vertex data.

Memory is coupled to the multiplication logic unit and the arithmetic logic unit. The memory has stored therein a plurality of constants and variables for being used in conjunction with the input buffer, the multiplication logic unit, the arithmetic logic unit, the register unit, the inverse logic unit, and the conversion module for processing the vertex data. Finally, an output converter is coupled to the output of

the arithmetic logic unit for being coupled to the lighting module to output the processed vertex data thereto.

To further assist integration, the lighting module includes a plurality of input buffers adapted for being coupled to a transform system for receiving vertex data therefrom. The input buffers include a first input buffer, a second input buffer, and a third input buffer. An input of the first buffer, the second input buffer, and the third input buffer are coupled to an output of the transform system.

10 Further included is a multiplication logic unit having a first input coupled to an output of the first input buffer and a second input coupled to an output of the second input buffer. An arithmetic logic unit has a first input coupled to an output of the second input buffer. The arithmetic logic unit further has a second input coupled to an output of the multiplication logic unit. An output of the arithmetic logic unit is coupled to the output of the lighting system.

20 Next provided is a first register unit having an input coupled to the output of the arithmetic logic unit and an output coupled to the first input of the arithmetic logic unit. A second register unit has an input coupled to the output of the arithmetic logic unit. Also, such second register has an output coupled to the first input and the second input of the multiplication logic unit. A lighting logic unit is also provided having a first input coupled to the output of the arithmetic logic unit, a second input coupled to the output of the first input buffer, and an output coupled to the first input of the multiplication logic unit.

25 Similar to the transform module, memory is coupled to at least one of the inputs of the multiplication logic unit and the output of the arithmetic logic unit. The memory has stored therein a plurality of constants and variables for being used in conjunction with the input buffers, the multiplication logic unit, the arithmetic logic unit, the first register unit, the second register unit, and the lighting logic unit for processing the vertex data.

Together, the foregoing transform/lighting architecture may work with a rasterizer that operates in homogeneous clip space to provide clip-less rasterization. This facilitates the placement of all of the components on the single semiconductor platform. In order to operate in homogeneous clip space, the rasterizer determines line equations for lines that define a primitive upon receipt of the primitive from an adjoining set-up module. Thereafter, a W-value is calculated using the line equations for points of intersections of the lines. An area is then determined based on the calculated W-values. Such area is representative of a portion of a display to be depicted. A space in the area is then identified using the line equations for rendering pixels therein.

These and other advantages of the present invention will become apparent upon reading the following detailed description and studying the various figures of the drawings.

2009-04-09 14:22:00

BRIEF DESCRIPTION OF THE DRAWINGS

The foregoing and other aspects and advantages are better understood from
5 the following detailed description of a preferred embodiment of the invention with
reference to the drawings, in which:

Figure 1 illustrates a prior art method of rasterization;

10 Figure 1A illustrates the association of a set of edge variables with a particular
pixel in accordance with a prior art method of rasterization;

15 Figure 1B is a flow diagram illustrating the various components of one
embodiment of the present invention implemented on a single semiconductor
platform;

Figure 2 is a schematic diagram of a vertex attribute buffer (VAB) in
accordance with one embodiment of the present invention;

20 Figure 2A is a chart illustrating the various commands that may be received
by VAB in accordance with one embodiment of the present invention;

25 Figure 2B is a flow chart illustrating a method of loading and draining vertex
attributes to and from VAB in accordance with one embodiment of the present
invention;

Figure 2C is a schematic diagram illustrating the architecture of the present
invention employed to implement the operations of Figure 2B;

30 Figure 3 illustrates the mode bits associated with VAB in accordance with
one embodiment of the present invention;

Figure 4 illustrates the transform module of the present invention;

5 Figure 4A is a flow chart illustrating a method of running multiple execution
threads in accordance with one embodiment of the present invention;

Figure 4B is a flow diagram illustrating a manner in which the method of
Figure 4A is carried out in accordance with one embodiment of the present
invention;

10 Figure 5 illustrates the functional units of the transform module of Figure 4
in accordance with one embodiment of the present invention;

15 Figure 6 is a schematic diagram of the multiplication logic unit (MLU) of the
transform module of Figure 5;

Figure 7 is a schematic diagram of the arithmetic logic unit (ALU) of the
transform module of Figure 5;

20 Figure 8 is a schematic diagram of the register file of the transform module
of Figure 5;

Figure 9 is a schematic diagram of the inverse logic unit (ILU) of the
transform module of Figure 5;

25 Figure 10 is a chart of the output addresses of output converter of the
transform module of Figure 5 in accordance with one embodiment of the present
invention;

Figure 11 is an illustration of the micro-code organization of the transform module of Figure 5 in accordance with one embodiment of the present invention;

5 Figure 12 is a schematic diagram of the sequencer of the transform module of Figure 5 in accordance with one embodiment of the present invention;

Figure 13 is a flowchart delineating the various operations associated with use of the sequencer of the transform module of Figure 12;

10 Figure 14 is a flow diagram delineating the operation of the sequencing component of the sequencer of the transform module of Figure 12;

15 Figure 14A is a flow diagram illustrating the components of the present invention employed for handling scalar and vector components during graphics-processing;

Figure 14B is a flow diagram illustrating one possible combination 1451 of the functional components of the present invention shown in Figure 14A which corresponds to the transform module of Figure 5;

20 Figure 14C is a flow diagram illustrating another possible combination 1453 of the functional components of the present invention shown in Figure 14A;

25 Figure 14D illustrates a method implemented by the transform module of Figure 12 for performing a blending operation during graphics-processing in accordance with one embodiment of the present invention;

Figure 15 is a schematic diagram of the lighting module of one embodiment of the present invention;

30

Figure 16 is a schematic diagram showing the functional units of the lighting module of Figure 15 in accordance with one embodiment of the present invention;

5 Figure 17 is a schematic diagram of the multiplication logic unit (MLU) of
the lighting module of Figure 16 in accordance with one embodiment of the present
invention;

10 Figure 18 is a schematic diagram of the arithmetic logic unit (ALU) of the
lighting module of Figure 16 in accordance with one embodiment of the present
invention;

Figure 19 is a schematic diagram of the register unit of the lighting module of
Figure 16 in accordance with one embodiment of the present invention;

15 Figure 20 is a schematic diagram of the lighting logic unit (LLU) of the
lighting module of Figure 16 in accordance with one embodiment of the present
invention;

20 Figure 21 is an illustration of the flag register associated with the lighting
module of Figure 16 in accordance with one embodiment of the present invention;

Figure 22 is an illustration of the micro-code fields associated with the
lighting module of Figure 16 in accordance with one embodiment of the present
invention;

25 Figure 23 is a schematic diagram of the sequencer associated with the lighting
module of Figure 16 in accordance with one embodiment of the present invention;

Figure 24 is a flowchart delineating the manner in which the sequencers of the transform and lighting modules are capable of controlling the input and output of the associated buffers in accordance with one embodiment of the present invention;

5 Figure 25 is a diagram illustrating the manner in which the sequencers of the transform and lighting modules are capable of controlling the input and output of the associated buffers in accordance with the method of Figure 24;

10 Figure 25B is a schematic diagram of the various modules of the rasterizer of
Figure 1B;

Figure 26 illustrates a schematic of the set-up module of the rasterization module of the present invention;

15 Figure 26A is an illustration showing the various parameters calculated by the set-up module of the rasterizer of Figure 26;

Figure 27 is a flowchart illustrating a method of the present invention associated with the set-up and traversal modules of the rasterizer component shown in Figure 26;

20 Figure 27A illustrates sense points that enclose a convex region that is moved to identify an area in a primitive in accordance with one embodiment of the present invention;

25 Figure 28 is a flowchart illustrating a process of the present invention associated with the process row operation 2706 of Figure 27;

Figure 28A is an illustration of the sequence in which the convex region of the present invention is moved about the primitive;

30

Figure 28B illustrates another example of the sequence in which the convex region of the present invention is moved about the primitive;

5 Figure 29 is a flowchart illustrating an alternate boustrophedonic process of the present invention associated with the process row operation 2706 of Figure 27;

Figure 29A is an illustration of the sequence in which the convex region of the present invention is moved about the primitive in accordance with the boustrophedonic process of Figure 29;

10

Figure 30 is a flowchart illustrating an alternate boustrophedonic process using boundaries;

15

Figure 31 is a flowchart showing the process associated with operation 3006 of Figure 30;

Figure 31A is an illustration of the sequence in which the convex region of the present invention is moved about the primitive in accordance with the boundary-based boustrophedonic process of Figures 30 and 31;

20

Figure 32 is a flowchart showing the process associated with operation 2702 of Figure 27;

25

Figure 32A is an illustration showing which area is drawn if no negative W-values are calculated in the process of Figure 32;

Figure 32B is an illustration showing which area is drawn if only one negative W-value is calculated in the process of Figure 32; and

Figure 32C is an illustration showing which area is drawn if only two negative W-values are calculated in the process of Figure 32.

DESCRIPTION OF THE PREFERRED EMBODIMENTS

Figures 1 and 1A show the prior art. Figures 1B - 32C show a graphics
5 pipeline system of the present invention.

Figure 1B is a flow diagram illustrating the various components of one embodiment of the present invention. As shown, the present invention is divided into four main modules including a vertex attribute buffer (VAB) 50, a transform module 52, a lighting module 54, and a rasterization module 56 with a set-up module 57. In one embodiment, each of the foregoing modules is situated on a single semiconductor platform in a manner that will be described hereinafter in greater detail. In the present description, the single semiconductor platform may refer to a sole unitary semiconductor-based integrated circuit or chip.
10

15

The VAB 50 is included for gathering and maintaining a plurality of vertex attribute states such as position, normal, colors, texture coordinates, etc. Completed vertices are processed by the transform module 52 and then sent to the lighting module 54. The transform module 52 generates vectors for the lighting module 54
20 to light. The output of the lighting module 54 is screen space data suitable for the set-up module which, in turn, sets up primitives. Thereafter, rasterization module 56 carries out rasterization of the primitives. It should be noted that the transform and lighting modules 52 and 54 might only stall on the command level such that a command is always finished once started.
25

30

In one embodiment, the present invention includes a hardware implementation that at least partially employs Open Graphics Library (OpenGL[®]) and D3DTM transform and lighting pipelines. OpenGL[®] is the computer industry's standard application program interface (API) for defining 2-D and 3-D graphic images. With OpenGL[®], an application can create the same effects in any operating system using any OpenGL[®]-adhering graphics adapter. OpenGL[®] specifies a set of

commands or immediately executed functions. Each command directs a drawing action or causes special effects.

Figure 2 is a schematic diagram of VAB 50 in accordance with one embodiment of the present invention. As shown, VAB 50 passes command bits 200 while storing data bits 204 representative of attributes of a vertex and mode bits 202. In use VAB 50 receives the data bits 204 of vertices and drains the same.

The VAB 50 is adapted for receiving and storing a plurality of possible vertex attribute states via the data bits 204. In use after such data bits 204, or vertex data, is received and stored in VAB 50, the vertex data is outputted from VAB 50 to a graphics-processing module, namely the transform module 52. Further, the command bits 200 are passed by VAB 50 for determining a manner in which the vertex data is inputted to VAB 50 in addition to other processing which will be described in greater detail with reference to Figure 2A. Such command bits 200 are received from a command bit source such as a microcontroller, CPU, data source or any other type of source which is capable of generating command bits 200.

Further, mode bits 202 are passed which are indicative of the status of a plurality of modes of process operations. As such, mode bits 202 are adapted for determining a manner in which the vertex data is processed in the subsequent graphics-processing modules. Such mode bits 202 are received from a command bit source such as a microcontroller, CPU, data source or any other type of source which is capable of generating mode bits 202.

25

It should be noted that the various functions associated with VAB 50 may be governed by way of dedicated hardware, software or any other type of logic. In various embodiments, 64, 128, 256 or any other number of mode bits 202 may be employed.

The VAB **50** also functions as a gathering point for the 64 bit data that needs to be converted into a 128-bit format. The VAB **50** input is 64 bits/cycle and the output is 128 bits/cycle. In other embodiments, VAB **50** may function as a gathering point for 128-bit data, and VAB **50** input may be 128 bits/cycle or any other combination. The VAB **50** further has reserved slots for a plurality of vertex attributes that are all IEEE 32 bit floats. The number of such slots may vary per the desires of the user. Table 1 illustrates exemplary vertex attributes employed by the present invention.

10

Table 1

15

Position: x,y,z,w
Diffuse Color: r,g,b,a
Specular Color: r,g,b
Fog: f
Texture0: s,t,r,q
Texture1: s,t,r,q
Normal: nx,ny,nz
Skin Weight: w

20

30

During operation, VAB **50** may operate assuming that the x,y data pair is written before the z,w data pair since this allows for defaulting the z,w pair to (0.0,1.0) at the time of the x,y write. This may be important for default components in OpenGL® and D3D™. It should be noted that the position, texture0, and texture1 slots default the third and fourth components to (0.0,1.0). Further, the diffuse color slot defaults the fourth component to (1.0) and the texture slots default the second component to (0.0).

The VAB **50** includes still another slot **205** used for assembling the data bits **204** that may be passed into or through the transform and lighting module **52** and **54**, respectively, without disturbing the data bits **204**. The data bits **204** in the slot **205** can be in a floating point or integer format. As mentioned earlier, the data bits **204**

of each vertex has an associated set of mode bits **202** representative of the modes affecting the processing of the data bits **204**. These mode bits **202** are passed with the data bits **204** through the transform and lighting modules **52** and **54**, respectively, for purposes that will be set forth hereinafter in greater detail.

5

In one embodiment, there may be 18 valid VAB, transform, and lighting commands received by VAB **50**. Figure **2A** is a chart illustrating the various commands that may be received by VAB **50** in accordance with one embodiment of the present invention. It should be understood that all load and read context commands, and the passthrough command shown in the chart of Figure **2A** transfer one data word of up to 128 bits or any other size.

10

Each command of Figure **2A** may contain control information dictating whether each set of data bits **204** is to be written into a high double word or low double word of one VAB address. In addition, a 2-bit write mask may be employed for providing control to the word level. Further, there may be a launch bit that informs VAB controller that all of the data bits **204** are present for a current command to be executed.

15

20

Each command has an associated stall field that allows a look-up to find information on whether the command is a read command in that it reads context memory or is a write command in that it writes context memory. By using the stall field of currently executing commands, the new command may be either held off in case of conflict or allowed to proceed.

25

In operation, VAB **50** can accept one input data word up to 128 bits (or any other size) per cycle and output one data word up to 128 bits (or any other size) per cycle. For the load commands, this means that it may take two cycles to load the data into VAB **50** to create a 128-bit quad-word and one cycle to drain it. For the scalar memories in the lighting module **54**, it is not necessary to accumulate a full quad-word, and these can be loaded in one cycle/address. For one vertex, it can take

up to 14 cycles to load the 7 VAB slots while it only takes 7 cycles to drain them. It should be noted, however, that it is only necessary to update the vertex state that changes between executing vertex commands. This means that, in one case, the vertex position may be updated taking 2 cycles, while the draining of the vertex data
5 takes 7 cycles. It should be noted that only 1 cycle may be required in the case of the x,y position.

Figure 2B is a flow chart illustrating one method of loading and draining
10 vertex attributes to and from VAB 50 during graphics-processing. Initially, in operation 210, at least one set of vertex attributes is received in VAB 50 for being processed. As mentioned earlier, each set of vertex attributes may be unique, and correspond to a single vertex.

15 In use the vertex attributes are stored in VAB 50 upon the receipt thereof in operation 212. Further, each set of stored vertex attributes is transferred to a corresponding one of a plurality of input buffers of the transform module 52. The received set of vertex attributes is also monitored in order to determine whether a received vertex attribute has a corresponding vertex attribute of a different set
20 currently stored in VAB 50, as indicated in operation 216.

Upon it being determined that a stored vertex attribute corresponds to the received vertex attribute in decision 217, the stored vertex attribute is outputted to the corresponding input buffer of the transform module 52 out of order. See
25 operation 218. Immediately upon the stored vertex attribute being outputted, the corresponding incoming vertex attribute may take its place in VAB 50. If no correspondence is found, however, each set of the stored vertex attributes may be transferred to the corresponding input buffer of the transform module 52 in accordance with a regular predetermined sequence. Note operation 219.

30

It should be noted that the stored vertex attribute might not be transferred in the aforementioned manner if it has an associated launch command. Further, in order for the foregoing method to work properly, the bandwidth of an output of VAB 50 must be at least the bandwidth of an input of VAB 50.

5

Figure 2C is a schematic diagram illustrating the architecture of the present invention employed to implement the operations of Figure 2B. As shown, VAB 50 has a write data terminal **WD**, a read data terminal **RD**, a write address terminal **WA**, and a read address **RA** terminal. The read data terminal is coupled to a first 10 clock-controlled buffer 230 for outputting the data bits 204 from VAB 50.

Also included is a first multiplexer 232 having an output coupled to the read address terminal of VAB 50 and a second clock-controlled buffer 234. A first input of the first multiplexer 232 is coupled to the write address terminal of VAB 50 while 15 a second input of the first multiplexer 232 is coupled to an output of a second multiplexer 236. A logic module 238 is coupled between the first and second multiplexers 232 and 236, the write address terminal of VAB 50, and an output of the second clock-controlled buffer 234.

20 In use the logic module 238 serves to determine whether an incoming vertex attribute is pending to drain in VAB 50. In one embodiment, this determination may be facilitated by monitoring a bit register that indicates whether a vertex attribute is pending or not. If it is determined that the incoming vertex attribute does have a match currently in VAB 50, the logic module 238 controls the first multiplexer 232 25 in order to drain the matching vertex attribute so that the incoming vertex attribute may be immediately stored in its place. On the other hand, if it is determined that the incoming vertex attribute does not have a match currently in VAB 50, the logic module 238 controls the first multiplexer 232 such that VAB 50 is drained and the incoming vertex attribute is loaded sequentially or in some other predetermined

order, per the input of the second multiplexer **236** which may be updated by the logic module **238**.

As a result, there is no requirement for VAB **50** to drain multiple vertex
5 attributes before a new incoming vertex attribute may be loaded. The pending vertex
attribute forces out the corresponding VAB counterpart if possible, thus allowing it
to proceed. As a result, VAB **50** can drain in an arbitrary order. Without this
capability, it would take 7 cycles to drain VAB **50** and possibly 14 more cycles to
load it. By overlapping the loading and draining, higher performance is achieved. It
10 should be noted that this is only possible if an input buffer is empty and VAB **50** can
drain into input buffers of the transform module **52**.

Figure 3 illustrates the mode bits associated with VAB **50** in accordance with
one embodiment of the present invention. The transform/light mode information is
15 stored in a register via mode bits **202**. Mode bits **202** are used to drive the
sequencers of the transform module **52** and lighting module **54** in a manner that will
be become apparent hereinafter. Each vertex has associated mode bits **202** that may
be unique, and can therefore execute a specifically tailored program sequence.
While, mode bits **202** may generally map directly to the graphics API, some of them
20 may be derived.

In one embodiment, the active light bits (LIS) of Figure 3 may be contiguous.
Further, the pass-through bit (VPAS) is unique in that when it is turned on, the
vertex data is passed through with scale and bias, and no transforms or lighting is
25 done. Possible mode bits **202** used when VPAS is true are the texture divide bits
(TDV0,1), and foggen bits (used to extract fog value in D3DTM). VPAS is thus used
for pre-transformed data, and TDV0,1 are used to deal with a cylindrical wrap mode
in the context of D3DTM.

Figure 4 illustrates the transform module of one embodiment of the present invention. As shown, the transform module 52 is connected to VAB 50 by way of 6 input buffers 400. In one embodiment, each input buffer 400 might be 7*128b in size. The 6 input buffers 400 each is capable of storing 7 quad words. Such input buffers 400 follow the same layout as VAB 50, except that the pass data is overlapped with the position data.

In one embodiment, a bit might be designated for each attribute of each input buffer 400 to indicate whether data has changed since the previous instance that the input buffer 400 was loaded. By this design, each input buffer 400 might be loaded only with changed data.

The transform module 52 is further connected to 6 output vertex buffers 402 in the lighting module 54. The output buffers include a first buffer 404, a second buffer 406, and a third buffer 408. As will become apparent hereinafter, the contents, i.e. position, texture coordinate data, etc., of the third buffer 408 are not used in the lighting module 54. The first buffer 404 and second buffer 406 are both, however, used for inputting lighting and color data to the lighting module 54. Two buffers are employed since the lighting module is adapted to handle two read inputs. It should be noted that the data might be arranged so as to avoid any problems with read conflicts, etc.

Further coupled to the transform module 52 is context memory 410 and micro-code ROM memory 412. The transform module 52 serves to convert object space vertex data into screen space, and to generate any vectors required by the lighting module 54. The transform module 52 also does processes skinning and texture coordinates. In one embodiment, the transform module 52 might be a 128-bit design processing 4 floats in parallel, and might be optimized for doing 4 term dot products.

30

Figure 4A is a flow chart illustrating a method of executing multiple threads in the transform module 52 in accordance with one embodiment of the present invention. In operation, the transform module 52 is capable of processing 3 vertices in parallel via interleaving. To this end, 3 commands can be simultaneously 5 executed in parallel unless there are stall conditions between the commands such as writing and subsequently reading from the context memory 410. The 3 execution threads are independent of each other and can be any command since all vertices contain unique corresponding mode bits 202.

10 As shown in Figure 4A, the method of executing multiple threads includes determining a current thread to be executed in operation 420. This determination might be made by identifying a number of cycles that a graphics-processing module requires for completion of an operation, and tracking the cycles. By tracking the cycles, each thread can be assigned to a cycle, thus allowing determination of the 15 current thread based on the current cycle. It should be noted, however, that such determination might be made in any desired manner that is deemed effective.

Next, in operation 422, an instruction associated with a thread to be executed during a current cycle is retrieved using a corresponding program counter number. 20 Thereafter, the instruction is executed on the graphics-processing module in operation 424.

In one example of use, the instant method includes first accessing a first instruction, or code segment, per a first program counter. As mentioned earlier, such 25 program counter is associated with a first execution thread. Next, the first code segment is executed in the graphics-processing module. As will soon become apparent, such graphics-processing module might take the form of an adder, a multiplier, or any other functional unit or combination thereof.

30 Since the graphics-processing module requires more than one clock cycle to complete the execution, a second code segment might be accessed per a second

program counter immediately one clock cycle after the execution of the first code segment. The second program counter is associated with a second execution thread, wherein each of the execution threads process a unique vertex.

5 To this end, the second code segment might begin execution in the graphics-processing module prior to the completion of the execution of the first code segment in the graphics-processing module. In use the graphics-processing module requires a predetermined number of cycles for every thread to generate an output. Thus, the various steps of the present example might be repeated for every predetermined
10 number of cycles.

This technique offers numerous advantages over the prior art. Of course, the functional units of the present invention are used more efficiently. Further, the governing code might be written more efficiently when the multiple threading
15 scheme is assumed to be used.

For example, in the case where the graphics-processing module includes a multiplier that requires three clock cycles to output an answer, it would be necessary to include two no operation commands between subsequent operations such as
20 $a=b*c$ and $d=e*a$, since “ a ” would not be available until after the three clock cycles. In the present embodiment, however, the code might simply call $d=e*a$ immediately subsequent $a=b*c$, because it can be assumed that such code will be executed as one of three execution threads that are called once every three clock cycles.

25 Figure 4B is a flow diagram illustrating a manner in which the method of Figure 4A is carried out. As shown, each execution thread has an associated program counter 450 that is used to access instructions, or code segments, in instruction memory 452. Such instructions might then be used to operate a graphics-processing module such as an adder 456, a multiplier 454, and/or an inverse logic
30 unit or register 459.

In order to accommodate a situation where at least two of the foregoing processing modules are used in tandem, at least one code segment delay **457** is employed between the graphics-processing modules. In the case where a three-thread framework is employed, a three-clock cycle code segment delay **457** is used.

5 In one embodiment, the code segment delay **457** is used when a multiplication instruction is followed by an addition instruction. In such case, the addition instruction is not executed until three clock cycles after the execution of the multiplication instruction in order to ensure that time has elapsed which is sufficient for the multiplier **456** to generate an output.

10

After the execution of each instruction, the program counter **450** of the current execution thread is updated and the program counter of the next execution thread is called by module **458** in a round robin sequence to access an associated instruction. It should be noted that the program counters might be used in any fashion including, but not limited to incrementing, jumping, calling and returning, performing a table jump, and/or dispatching. Dispatching refers to determining a starting point of code segment execution based on a received parameter. Further, it is important to understand that the principles associated with the present multiple thread execution framework might also be applied to the lighting module **54** of the graphics-processing pipeline of the present invention.

20

In the case where a three-thread framework is employed, each thread is allocated one input buffer and one output buffer at any one time. This allows loading of three more commands with data while processing three commands. The input buffers and output buffers are assigned in a round robin sequence in a manner that will be discussed later with reference to Figures **27** and **28**.

25

The execution threads are thus temporally and functionally interleaved. This means that each function unit is pipelined into three stages and each thread occupies one stage at any one time. In one embodiment, the three-threads might be set to always execute in the same sequence, i.e. zero then one then three. Conceptually,

the threads enter a function unit at t = clock modulo three. Once a function unit starts work, it takes three cycles to deliver the result (except the ILU that takes six), at which time the same thread is again active.

5 Figure 5 illustrates the functional units of the transform module 52 of Figure 4 in accordance with one embodiment of the present invention. As shown, included are input buffers 400 that are adapted for being coupled to VAB 50 for receiving vertex data therefrom.

10 A memory logic unit (MLU) 500 has a first input coupled to an output of input buffers 400. As an option, the output of MLU 500 might have a feedback loop 502 coupled to the first input thereof.

15 Also provided is an arithmetic logic unit (ALU) 504 having a first input coupled to an output of MLU 500. The output of ALU 504 further has a feedback loop 506 connected to the second input thereof. Such feedback loop 502 may further have a delay 508 coupled thereto. Coupled to an output of ALU 504 is an input of a register unit 510. It should be noted that the output of register unit 510 is coupled to the first and second inputs of MLU 500.

20 An inverse logic unit (ILU) 512 is provided including an input coupled to the output of ALU 504 for performing an inverse or an inverse square root operation. In an alternate embodiment, ILU 512 might include an input coupled to the output of register unit 510.

25 Further included is a conversion, or smearing, module 514 coupled between an output of ILU 512 and a second input of MLU 500. In use the conversion module 514 serves to convert scalar vertex data to vector vertex data. This is accomplished by multiplying the scalar data by a vector so that the vector operators such as the
30 multiplier and/or adder may process it. For example, a scalar A, after conversion,

may become a vector (A,A,A,A). In an alternate embodiment, the smearing module **514** might be incorporated into the multiplexers associated with **MLU 500**, or any other component of the present invention. As an option, a register **516** might be coupled between the output of **ILU 512** and an input of the conversion unit **514**.

5 Further, such register **516** might be threaded.

Memory **410** is coupled to the second input of **MLU 500** and the output of **ALU 504**. In particular, memory **410** has a read terminal coupled to the second input of **MLU 500**. Further, memory **410** has a write terminal coupled to the output of

10 **ALU 504**.

The memory **410** has stored therein a plurality of constants and variables for being used in conjunction with the input buffer **400**, **MLU 500**, **ALU 504**, register unit **510**, **ILU 512**, and the conversion module **514** for processing the vertex data.

15 Such processing might include transforming object space vertex data into screen space vertex data, generating vectors, etc.

Finally, an output converter **518** is coupled to the output of **ALU 504**. The output converter **518** serves for being coupled to a lighting module **54** via output buffers **402** to output the processed vertex data thereto. All data paths except for the **ILU** might be designed to be 128 bits wide or other data path widths may be used.

Figure **6** is a schematic diagram of **MLU 500** of the transform module **52** of Figure **5** in accordance with one embodiment of the present invention. As shown, **MLU 500** of the transform module **52** includes four multipliers **600** that are coupled in parallel.

MLU 500 of transform module **52** is capable of multiplying two four component vectors in three different ways, or pass one four component vector.

MLU 500 is capable of performing multiple operations. Table 2 illustrates such operations associated with MLU 500 of transform module 52.

Table 2

5	CMLU_MULT o[0] = a[0]*b[0],o[1] = a[1]*b[1],o[2] = a[2]*b[2],o[3] = a[3]*b[3]
	CMLU_MULA o[0] = a[0]*b[0],o[1] = a[1]*b[1],o[2] = a[2]*b[2],o[3] = a[3]
	CMLU_MULB o[0] = a[0]*b[0],o[1] = a[1]*b[1],o[2] = a[2]*b[2],o[3] = b[3]
10	CMLU_PASA o[0] = a[0],o[1] = a[1],o[2] = a[2],o[3] = a[3]
	CMLU_PASB o[0] = b[0],o[1] = b[1],o[2] = b[2],o[3] = b[3]

Possible A and B inputs are shown in Table 3.

Table 3

15	MA_M	MLU
	MA_V	Input Buffer
	MA_R	RLU (shared with MB_R)
20	MB_I	ILU
	MB_C	Context Memory
	MB_R	RLU (shared with MA_R)

Table 4 illustrates a vector rotate option capable of being used for cross products.

Table 4

30	MR_NONE	No change
	MR_ALBR	Rotate A[XYZ] vector left, B[XYZ] vector right
	MR_ARBL	Rotate A[XYZ] vector right, B[XYZ] vector left

Figure 7 is a schematic diagram of ALU 504 of transform module 52 of Figure 5 in accordance with one embodiment of the present invention. As shown, 35 ALU 504 of transform module 52 includes three adders 700 coupled in parallel/series. In use ALU 504 of transform module 52 can add two three component vectors, pass one four component vector, or smear a vector component

across the output. Table 5 illustrates various operations of which ALU **504** of transform module **52** is capable.

Table 5

5	CALU_ADDA	$o[0] = a[0]+b[0], o[1] = a[1]+b[1], o[2] = a[2]+b[2], o[3] = a[3]$
	CALU_ADDB	$o[0] = a[0]+b[0], o[1] = a[1]+b[1], o[2] = a[2]+b[2], o[3] = b[3]$
	CALU_SUM3B	$o[0123] = b[0] + b[1] + b[2]$
10	CALU_SUM4B	$o[0123] = b[0] + b[1] + b[2] + b[3]$
	CALU_SMRB0	$o[0123] = b[0]$
	CALU_SMRB1	$o[0123] = b[1]$
	CALU_SMRB2	$o[0123] = b[2]$
	CALU_SMRB3	$o[0123] = b[3]$
15	CALU_PASA	$o[0] = a[0], o[1] = a[1], o[2] = a[2], o[3] = a[3]$
	CALU_PASB	$o[0] = b[0], o[1] = b[1], o[2] = b[2], o[3] = b[3]$

Table 6 illustrates the A and B inputs of ALU **504** of transform module **52**.

Table 6

20	AA_A	ALU (one instruction delay)
	AA_C	Context Memory
	AB_M	MLU

25 It is also possible to modify the sign bits of the A and B input by effecting no change, negation of B, negation of A, absolute value A,B. It should be noted that when ALU **504** outputs scalar vertex data, this scalar vertex data is smeared across the output in the sense that each output represents the scalar vertex data. The pass control signals of MLU **500** and ALU **504** are each capable of disabling all special 30 value handling during operation.

Figure 8 is a schematic diagram of the vector register file **510** of transform module **52** of Figure 5 in accordance with one embodiment of the present invention. As shown, the vector register file **510** includes four sets of registers **800** each having 35 an output connected to a first input of a corresponding multiplexer **802** and an input coupled to a second input of the corresponding multiplexer **802**.

In one embodiment of the present invention, the vector register file **510** is threaded. That is, there are three copies of the vector register file **510** and each thread has its own copy. In one embodiment, each copy contains eight registers, each of which might be 128 bits in size and store four floats. The vector register file **510** is written from ALU **504** and the output is fed back to MLU **500**. The vector register file **510** has one write and one read per cycle.

In operation, it is also possible to individually mask a write operation to each register component. The vector register file **510** exhibits zero latency when the write address is the same as the read address due to a bypass path **511** from the input to the output. In this case, unmasked components would be taken from the registers and masked components would be bypassed. The vector register file **510** is thus very useful for building up vectors component by component, or for changing the order of vector components in conjunction with the ALU SMR operations (See Table 5).

Temporary results might be also stored in the vector register file **510**.

Figure **9** is a schematic diagram of ILU **512** of transform module **52** of Figure **5** in accordance with one embodiment of the present invention. As shown, ILU **512** of transform module **52** is capable of generating a floating-point reciprocal ($1/D$) and a reciprocal square root ($1/D^{(1/2)}$). To carry out such operations, either one of two iterative processes might be executed on a mantissa. Such processes might be executed with any desired dedicated hardware, and are shown below:

	<u>Reciprocal (1/D)</u>	<u>Reciprocal Square-root (1/D^(1/2))</u>
25	$x_n + 1 = x_n(2 - x_n * D)$	$x_{n+1} = (1/2) * x_n(3 - x_n^2 * D)$
	1) table look up for x_n (seed) x_n	table look up for x_n (seed) $x_n * x_n$
30	2) 1 st iteration: multiply-add	1 st iteration: multiply-add

	$2-x_n*D$	$3-x_n^2*D$
3)	1 st iteration: multiply $x_n(2-x_n*D)$	1 st iteration: multiply $(1/2)*x_n(3-x_n^2*D)$
5		
4)	2 nd iteration: no-op pass x_{n+1}	2 nd iteration: square x_{n+1}^2
10		
5)	2 nd iteration: multiply-add $2-x_{n+1}*D$	2 nd iteration: multiply-add $3-x_{n+1}^2*D$
15		
6)	2 nd iteration: multiply $x_{n+1}(2-x_{n+1}*D)$	2 nd iteration: multiply $(1/2)*x_{n+1}(3-x_{n+1}^2*D)$

15 As shown, the two processes are similar, affording a straightforward design.
It should be noted that the iterations might be repeated until a threshold precision is met.

20 In operation, ILU **512** performs two basic operations including an inverse operation and inverse square root operation. Unlike the other units, it requires six cycles to generate the output. The input is a scalar, and so is the output. As set forth earlier, the threaded holding register **516** at ILU **512** output is relied upon to latch the result until the next time a valid result is generated. Further, the scalar output is smeared into a vector before being fed into MLU **500**. The inverse unit **512** uses
25 look-up tables and a two pass Newton-Raphson iteration to generate IEEE (Institute of Electrical and Electronics Engineers) outputs accurate to within about 22 mantissa bits. Table 7 illustrates the various operations that might be performed by ILU **512** of transform module **52**.

CILU_INV	$o = 1.0/a$
CILU_ISQ	$o = 1.0/\sqrt{a}$
CILU_CINV	$o = 1.0/a$ (with range clamp)
CILU_NOP	no output

5

The foregoing range clamp inversion operation of Table 7 might be used to
allow clipping operations to be handled by rasterization module 56. Coordinates are
transformed directly into screen space that can result in problems when the
homogeneous clip space w is near 0.0. To avoid multiplying by 1.0/0.0 in the
10 perspective divide, the 1/w calculation is clamped to a minimum and a maximum
exponent.

In use the context memory 410 as shown in Figure 5 reads and writes only
using quad-words. The memory can be read by MLU 500 or ALU 504 each cycle,
15 and can be written by ALU 504. Only one memory read is allowed per cycle. If a
read is necessary, it is done at the start of an instruction and then pipelined down to
ALU 504 three cycles later. Context memory 410 need not necessarily be threaded.

Figure 10 is a chart of the output addresses of output converter 518 of
20 transform module 52 of Figure 5 in accordance with one embodiment of the present
invention. The output converter 518 is responsible for directing the outputs to
proper destinations, changing the bit precision of data, and some data swizzling to
increase performance. All data destined for lighting module 54 is rounded to a 22
bit floating point format organized as S1E8M13 (one sign, eight exponent, 13
25 mantissa bits). The destination buffers 402 as shown in Figure 4 in lighting module
54 are threaded.

Data swizzling is useful when generating vectors. Such technique allows the
generation of a distance vector (1,d,d*d) without penalty when producing a vector.
30 The distance vector is used for fog, point parameter and light attenuation. This is
done with an eye vector and light direction vectors. Table 8 illustrates the various
operations associated with such vectors. It should be noted that, in the following

table, squaring the vector refers to $d^2 = \text{dot}[(x,y,z), (x,y,z)]$, and storing d^2 in the w component of (x,y,z) .

Table 8

5

1. Square the vector (x,y,z,d*d) (output d*d to VBUF, 1.0 to VBUF)
2. Generate inverse sqrt of d*d (1/d)
3. Normalize vector (x/d,y/d,z/d,d) (output x/d,y/d,z/d to WBUF, d to VBUF)

10

It should be noted that the math carried out in the present invention might not always be IEEE compliant. For example, it might be assumed that “0” multiplied by any number renders “0.” This is particularly beneficial when dealing with the equations such as $d=d^2*1/(d^2)^{1/2}$, where $d=0$. Without making the foregoing assumption, such equation would afford an error, thus causing problems in making related computations.

Figure 11 is an illustration of the micro-code organization of transform module 52 of Figure 5 in accordance with one embodiment of the present invention.

20

The transform module micro-code might be arranged into 15 fields making up a total width of 44 bits. Fields might be delayed to match the data flow of the units. MLU 500 operations are executed at a delay of zero, ALU operations are executed at a delay of one, and RLU, output operations are executed at a delay of two. Each delay is equivalent to three cycles.

25

Figure 12 is a schematic diagram of sequencer 1200 of transform module 52 of Figure 5 in accordance with one embodiment of the present invention. As shown in Figure 12, sequencer 1200 of transform module 52 includes a buffer 1202 adapted for receiving the mode bits from VAB 50 that are indicative of the status of a plurality of modes of process operations.

Also included is memory **412** capable of storing code segments that each are adapted to carry out the process operations in accordance with the status of the modes. A sequencing module **1206** is coupled between memory **412** and a control vector module **1205** which is in turn coupled to buffer **1202** for identifying a plurality of addresses in memory **412** based on a control vector derived from mode bits **202**. The sequencing module **1206** is further adapted for accessing the addresses in memory **412** for retrieving the code segments that might be used to operate transform module **52** to transfer data to an output buffer **1207**.

Figure 13 is a flowchart delineating the various operations associated with use of sequencer **1200** of transform module **52** of Figure 12. As shown, sequencer **1200** is adapted for sequencing graphics-processing in a transform or lighting operation. In operation **1320**, mode bits **202** are first received which are indicative of the status of a plurality of modes of process operations. In one embodiment, mode bits **202** might be received from a software driver.

Then, in operation **1322**, pluralities of addresses are then identified in memory based on mode bits **202**. Such addresses are then accessed in the memory in operation **1324** for retrieving code segments that each are adapted to carry out the process operations in accordance with the status of the modes. The code segments are subsequently executed with a transform or lighting module for processing vertex data. Note operation **1326**.

Figure 14 is a flow diagram delineating the operation of the sequencing module **1206** of sequencer **1200** of transform module **52** of Figure 12. As shown, a plurality of mode registers **1430** each include a unique set of mode bits **202** which in turn correspond to a single vertex. It should be noted that mode registers **1430** are polled in a round robin sequence in order to allow the execution of multiple execution threads in the manner set forth earlier during reference to Figures 4A and 4B.

Once the current execution thread is selected, a corresponding group of mode bits **202** are decoded in operation **1432**. Upon mode bits **202** being decoded in operation **1432**, a control vector is afforded which includes a plurality of bits each of which indicate whether a particular code segment is to be accessed in ROM **1404** for processing the corresponding vertex data.

Upon determining whether a code segment should be accessed in ROM **1404** and executed, a pointer operation **1436** increments the current thread pointer to start the next execution thread to obtain a second group mode bits **202** to continue a similar operation. This might be continued for each of the threads in a round robin sequence.

Once the control vector has been formed for a particular group of mode bits **202**, a priority encoder operation **1438** determines, or identifies, a next “1” or enabled, bit of the control vector. If such a bit is found, the priority encoder operation **1438** produces an address in ROM **1404** corresponding to the enabled bit of the control vector for execution purposes.

Upon returning to the initial group of mode bits **202** after handling the remaining threads, and after the mode bits have been decoded and the control vector is again available, a masking operation **1434** might be used to mask the previous “1”, or enabled, bit that was identified earlier. This allows analysis of all remaining bits after mask operation **1434**.

25

The foregoing process might be illustrated using the following tables. Table 9 shows a plurality of equations that might be executed on subject vertex data.

Table 9

$$\begin{aligned} R &= (a) \\ R &= (a + d * e) \\ R &= (a + b * c + f) \end{aligned}$$

30

5

R = (a + b*c + d*e)
R = 1.0/(a)
R = 1.0/(a + d*e)
R = 1.0/(a + b*c + f)
R = 1.0/(a + b*c + d*e)

As shown, there are four possibilities of products that might be summed in addition to an inverse operation (a, b*c, d*e, f, and 1/x). Next, mode fields might be
10 defined. Table 10 illustrates a pair of mode fields, mode.y and mode.z, each having assigned thereto a predetermined set of the operations of Table 9.

Table 10

15 mode.y[4] 0: R = a
1: R = a + d*e
2: R = a + b*c + f
3: R = a + b*c + d*e
20 mode.z[2] 0: R = R
1: R = 1.0/R

Thereafter, each of the operations might be positioned in memory with an associated address. Table 11 illustrates a plurality of memory addresses each having
25 an associated operation. Also shown is a set of control vector definitions.

Table 11

30 ROM[0]: R = a
ROM[1]: R = R + b*c
ROM[2]: R = R + d*e
ROM[3]: R = R + f
ROM[4]: R = 1.0/R
35 cv[0] = 1;
cv[1] = (mode.y==2 || mode.y==3) ? 1 : 0;
cv[2] = (mode.y==1 || mode.y==3) ? 1 : 0;
cv[3] = (mode.y==2) ? 1 : 0;
cv[4] = (mode.z==1) ? 1 : 0;
40

Table 12 illustrates the execution of an example.

Table 12

45

R = a+d*e corresponds to:

mode.y = 1;
mode.z = 0;

which in turn affords the following control vector:

5

cv[0] = 1;
cv[1] = 0;
cv[2] = 1;
cv[3] = 0;
cv[4] = 0;

10

execution
first cycle:

15

cv[0] is TRUE so execute ROM[0]
more TRUE values in control vector, so do not terminate
program

20

second cycle:

25

cv[1] is FALSE so keep looking
cv[2] is TRUE so execute ROM[2]
no more TRUE values in control vector, so terminate
program

30

As such, sequencer **1200** of transform module **52** steps through a threaded control vector which is derived from threaded mode bits **202**, and executes every ROM address whose corresponding control vector bit is set to “TRUE”. The control vector has the same length as the ROM. The sequencer **1200** is capable of stepping through an arbitrary control vector at the rate of one “1”, or enabled bit per a predetermined number of cycles. Commands that do not use mode bits **202** might be executed by on-the-fly micro-code generation due to the simplicity thereof.

35

By representing such statuses by way of a unique string of mode bits **202**, it is unnecessary to execute a plurality of if-then clauses in the graphics-processing hardware to determine the statuses of the various operations. Improved performance is thereby afforded. Conceptually, it is as if the if clauses in a program language had been moved to sequencer **1200** which in turn instantly skips instructions with a “FALSE” condition, as indicated by mode bits **202**.

40

As indicated earlier, code segments are stored in the ROM which are capable of handling the various statuses of the operations identified by the mode bits. In one embodiment a separate code segment might be retrieved for handling each operation indicated by the mode bits. In the alternative, a single comprehensive code segment

5 might be written for handling each or some combinations of operations that are possible. It should be noted, however, that generating such large code segments for each combination of operations requires additional code space, and it therefore might be beneficial to modularize the code segments for only commonly used combinations of operations.

10

Since mode bits **202** do not change once the vertex commences execution, the control vector generation might only have to be done once per vertex before entering the sequencer. Exceptions to this might arise in some cases, however, such as lighting where operations might be repeated. When the last vertex instruction is

15

found, an end of sequence (EOS) signal might be asserted. This in turn might be used to change the status of the input and output buffers, and to allow the start of the next command in a manner that will be set forth during reference to Figures **28A** and **28B**. It should be noted that the EOS signal is pipeline delayed for release of the destination buffer similar to the manner in which the instructions are handled. See

20

Figure **4B**.

Figure **14A** is a flow diagram illustrating the various functional components of the present invention employed for integrating the handling of scalar and vector vertex data during graphics-processing. As shown, one functional aspect **1440**

25

includes inputting vector vertex data into a processing module, i.e. adder, multiplier, etc., for outputting vector vertex data. In another functional aspect **1442**, vector vertex data is processed by a vector processing module, i.e. adder, multiplier, etc., which outputs scalar vertex data that is in turn converted, or smeared, again into vector vertex data.

30

In yet another functional aspect 1444, vector vertex data is masked, thereby converted to scalar vertex data, after which it is stored in memory, i.e. register logic unit, for the purpose of generating vector vertex data. In still yet another functional aspect 1446, scalar vertex data is extracted by a vector processing module, i.e. adder, multiplier, etc., which in turn is processed by a scalar processing module, i.e. inverse logic unit, which renders scalar vertex data. This scalar vertex data is converted again into vector vertex data.

Figure 14B is a flow diagram illustrating one possible combination 1451 of the functional components of the present invention shown in Figure 14A which corresponds to transform module 52 of Figure 5. It should be noted that functional aspects 1444 and 1446 might have delays associated therewith in a manner similar to that set forth earlier during reference to Figure 4B. Figure 14C is a flow diagram illustrating yet another possible combination 1453 of the functional components of the present invention shown in Figure 14A.

Multiplexers might accomplish the extraction of the scalar vertex data from the vector vertex data in the functional modules of Figures 14A-14C. Such multiplexers might also be responsible for any data swizzling that might be required before processing by the various functional modules. In one embodiment, the multiplexers might be capable of passing and rotating vector vertex data, and rely on other graphics-processing modules such as an ALU for other processing. In yet another embodiment, the multiplexers might be capable of arbitrarily rearranging attributes independently without penalty.

25

Figure 14D illustrates a method in which the transform system is adapted for performing a blending, or skinning operation during graphics-processing in a graphics pipeline via a hardware implementation such as an application specific integrated circuit (ASIC). During processing in the pipeline, in operation 1470, a plurality of matrices, a plurality of weight values each corresponding with one of the

matrices, and vertex data are received. It should be noted that an additional set of matrices might be required for normal vertex data.

Subsequently, in operation **1472**, a sum of a plurality of products is then
5 calculated with each product being calculated by the multiplication of the vertex
data, one of the matrices and the weight corresponding to the matrix. Such sum of
products is then outputted in operation **1474** for additional processing.

In summary, the following sum of products might be calculated:

10

Equation #1

$$v' = \sum w_i * M_i * v \text{ for } i=1 \dots x$$

where v = inputted vertex data

w = weight value

15

M = matrix

x = number of matrices

v' = vertex data for output to a processing module

Equation #2

20

$$n' = \sum w_i * I_i * n \text{ for } i=1 \dots x$$

where n = inputted vertex data (normal vector)

w = weight value

I = inverted matrix (inverse transpose matrix)

x = number of inverted matrices

25

n' = vertex data for output to a processing
module (normal vector)

Equation #3

30

$$v_s = [O_x, O_y, O_z, \emptyset]^T$$

$$1/(v''_{wc}) * [(v''_x), (v''_y), (v''_z), 1]'$$

where $v'' = C * v'$

$v' = \text{sum of products from Equation } \#1$

$C = [S_x, S_y, S_z, 1]' * P$

5 $P = \text{projection matrix}$

$v_s = \text{screen vector for display purposes}$

$O = \text{viewport offset}$

$S = \text{viewport scale}$

10 It should be noted that there are many ways to represent the weights w_i set forth hereinabove. For example, in Equations #1 and #2 above, it might be said that $i = 1 \dots (x-1)$, leaving $w_x (w_i \text{ where } i = x)$ to be calculated by the equation $1 - \sum w_i$. By representing the weights w_i in this way, it is ensured that all of the weights w sum to 1.

15 In one embodiment, the matrices might include model view matrices (M), and the sum of products (v') might be outputted for additional processing by a lighting operation. See Equation #1. This sum of products (v') might also be used to generate another sum of products (v_s) for display purposes by using a composite matrix (C). See Equation #3. Still yet, the matrices might include inverse matrices (I) and the vertex data might include normal vector data (n). In such case, the additional processing might include a lighting operation. See Equation #2.

25 Figure 15 is a schematic diagram of lighting module 54 in accordance with one embodiment of the present invention. As shown, lighting module 54 includes buffers 402 to which transform module 52 outputs the vertex data. As shown, buffer 408 bypasses lighting module 54 by way of the pathway 1501. Further coupled to lighting module 54 is a context memory 1500 and micro-code ROM memory 1502.

30 The lighting module 54 is adapted for handling lighting in addition to fog and point parameters. In use lighting module 54 controls the buffer bypass pathway

1501, and calculates the diffuse, point size, and specular output colors as well as the fog value. It should be noted that lighting module 54 employs the same mode bits 202 as transform module 52.

5 The lighting module 54 further requires less precision with respect to transform module 52, and therefore processes 22 bit floating point values (1.8.13 format) organized in tri-words. Since the data of third buffer 408 is 128 bits, it utilizes bypass pathway 1501 around lighting module 54. The lighting module 54 is event driven and simultaneously executes three threads in a manner similar to
10 transform module 52 as was set forth earlier with reference to Figures 4A and 4B. It should be noted that lighting module 54 might require command launch approval from an outside source.

15 Figure 16 is a schematic diagram showing the functional units of lighting module 54 of Figure 15 in accordance with one embodiment of the present invention. As shown, included are input buffers 402 adapted for being coupled to a transform system for receiving vertex data therefrom. As set forth earlier, input buffers 402 include a first input buffer 404, a second input 406, and a third input buffer 408. An input of first buffer 404, second input buffer 406, and third input buffer 408 are coupled to an output of transform module 52. For bypass purposes, the output of third buffer 408 is coupled to the output of lighting module 54 via a delay 1608.
20
25

Further included is a MLU 1610 having a first input coupled to an output of first input buffer 404 and a second input coupled to an output of second input buffer 406. The output of MLU 1610 has a feedback loop 1612 coupled to the second input thereof. An arithmetic logic unit (ALU) 1614 has a first input coupled to an output of second input buffer 406. ALU 1614 further has a second input coupled to an output of MLU 1610. An output of ALU 1614 is coupled to the output of lighting module 54. It should be noted that the output of ALU 1614 and the output of the
30

third input buffer **408** are coupled to the output of lighting module **54** by way of multiplexer **1616**.

Next provided is a first register unit **1618** having an input coupled to the
5 output of ALU **1614** and an output coupled to the first input of ALU **1614**. A second register unit **1620** has an input coupled to the output of ALU **1614**. Also, such second register **1620** has an output coupled to the first input and the second input of MLU **1610**.

10 A lighting logic unit (LLU) **1622** is also provided having a first input coupled to the output of ALU **1614**, a second input coupled to the output of the first input buffer **404**, and an output coupled to the first input of MLU **1610**. It should be noted that the second input of LLU **1622** is coupled to the output of the first input buffer **404** via a delay **1624**. Further, the output of LLU **1622** is coupled to the first input
15 of MLU **1610** via a first-in first-out register unit **1626**. As shown in Figure **16**, the output of LLU **1622** is also coupled to the first input of MLU **1610** via a conversion module **1628**. In operation, such conversion module **1628** is adapted for converting scalar vertex data to vector vertex data in a manner similar to that of transform module **52**.

20 Finally, memory **1500** is coupled to at least one of the inputs of MLU **1610** and the output of arithmetic logic unit **1614**. In particular, memory **1610** has a read terminal coupled to the first and the second input of MLU **1610**. Further, memory **1500** has a write terminal coupled to the output of ALU **1614**.

25 The memory has stored therein a plurality of constants and variables for being used in conjunction with input buffers **402**, MLU **1610**, ALU **1614**, first register unit **1618**, second register unit **1620**, and LLU **1622** for processing the vertex data.

30

Figure 17 is a schematic diagram of MLU 1610 of lighting module 54 of Figure 16 in accordance with one embodiment of the present invention. As shown, MLU 1610 of lighting module 54 includes three multipliers 1700 in parallel. In operation, the present MLU 1610 is adapted to multiply two three component vectors, or pass one three component vector. The multiplication of the three component vectors might be accomplished by way of a dot product or a parallel multiply. Table 13 illustrates the operations that MLU 1610 of lighting module 54 is capable of performing.

10

Table 13

ZMLU_MULT	$o[0] = a[0]*b[0]$, $o[1] = a[1]*b[1]$, $o[2] = a[2]*b[2]$
ZMLU_PASA	$o[0] = a[0]$, $o[1] = a[1]$, $o[2] = a[2]$
ZMLU_PASB	$o[0] = b[0]$, $o[1] = b[1]$, $o[2] = b[2]$

15

Table 14 illustrates the possible A and B inputs of MLU 1610 of lighting module 54.

20

Table 14

25

MA_V	VBUFFER
MA_L	LLU
MA_R	RLU[2,3] (shared with MB_R)
MA_C	Context memory (shared with MB_C)
MB_M	MLU
MB_W	WBUFFER
MB_R	RLU[2,3] (shared with MA_R)
MB_C	Context memory (shared with MA_C)

30

Figure 18 is a schematic diagram of ALU 1614 of lighting module 54 of Figure 16 in accordance with one embodiment of the present invention. As shown, ALU 1614 includes three adders 1800 in parallel/series. In use ALU 1614 is capable of adding two three component vectors, or passing one three component vector.

Table 15 illustrates the various operations of which ALU **1614** of lighting module **54** is capable.

Table 15

5

ZALU_ADD	$o[0] = a[0]+b[0]$, $o[1] = a[1]+b[1]$, $o[2] = a[2]+b[2]$
ZALU_SUM3B	$o[012] = b[0] + b[1] + b[2]$
ZALU_PASA	$o[0] = a[0]$, $o[1] = a[1]$, $o[2] = a[2]$
ZALU_PASB	$o[0] = b[0]$, $o[1] = b[1]$, $o[2] = b[2]$

10

Table 16 illustrates the possible A and B inputs to ALU **1614** of lighting module **54**.

Table 16

15

AA_W	WBUFFER
AA_R	RLU[0,1]
AB_M	MLU

20

Figure 19 is a schematic diagram of register units **1618** and **1620** of lighting module **54** of Figure 16 in accordance with one embodiment of the present invention. As shown, register units **1618** and **1620** each include two sets of registers **1900** each having an output connected to a first input of a corresponding multiplexer **1902** and an input coupled to a second input of multiplexer **1902**.

25

Register units **1618** and **1620** of lighting module **54** are split into two registers for ALU **1614** and two registers for MLU **1610**. In one embodiment, the registers are threaded. The register units **1618** and **1620** exhibit zero latency when a write address is the same as a read address due to a bypass path from the input to the outputs.

Figure 20 is a schematic diagram of LLU **1622** of lighting module **54** of Figure 16 in accordance with one embodiment of the present invention. LLU **1622** is the lighting unit of lighting module **54**. It is a scalar block that computes lighting

coefficients later used to multiply the light+material colors. LLU **1622** includes two MAC's, an inverter, four small memories, and a flag register.

The flag register is used to implement the conditional parts of the lighting
5 equations. The outputs are an ambient, diffuse, and specular coefficient. The scalar
memories contain variables used for the specular approximations and constants. The
first location of each memory contains 1.0 (for ctx0 and ctx2) and 0.0 (for ctx1 and
ctx3). In one embodiment, these are hardwired and do not need to be loaded.

10 In use LLU **1622** fundamentally implements the equation: $(x + L) / (M*x+N)$. This equation is used to approximate a specular lighting term. The inputs to LLU **1622** are from ALU **1614** of lighting module **54** and are the dot products used in the lighting equations. As set forth earlier, with respect to Figure 16, there is an output FIFO **1626** between LLU **1622** and MLU **1610** which buffers
15 coefficients until MLU **1610** needs them. In one embodiment, such FIFO **1626** might be threaded along with delays **1608** and **1624**, and registers **1618** and **1620**. Due to possible color material processing, it is unknown when the diffuse and specular outputs are consumed by MLU **1610**.

20 There is specially adapted hardware for dealing with the diffuse output alpha component since lighting module **54** only deals with R,G,B components. Such specially adapted hardware is capable of outputting two types of alpha components, namely vtx color α [Tbuffer], and stored ctx α [Ctx store]. The choice between the foregoing alpha components is governed by mode bits **202**.

25 In operation, LLU **1622** calculates ambient (Ca), diffuse (Cde), and specular (Cs) coefficients of lighting. These coefficients are then multiplied with the ambient, diffuse, and specular colors to generate a light's contribution to the vertex color.
Table 16A includes a list of inputs received by LLU **1622** and the calculations
30 carried out to generate the ambient (Ca), diffuse (Cde), and specular (Cs) coefficients of lighting. It should be noted that any desired hardware configuration

might be employed to implement LLU 1622. In one embodiment, the specific configuration shown in Figure 20 might be employed.

Table 16A

5

Input definitions:

10	n = normal vector	(from transform engine)
	e = normalized eye vector	(from transform engine)
	l = normalized light vector	(from transform engine)
	s = spotlight vector*light vector	(from transform engine)
	D = distance vector $(1,d,d^2)$	(from transform engine)
	h = half angle vector	(from lighting engine)
	K = attenuation constant vector (K_0, K_1, K_2)	(from context memory)

The LLU might receive the following scalar data in carrying out its calculations:

20	n*l (from MLU/ALU)
	n*h (from MLU/ALU)
	K*D (from MLU/ALU)
	s (from transform engine)
	power0 (material exponent from ctx0-3 memory)
	power1 (spotlight exponent from ctx0-3 memory)
25	range (from ctx0-3 memory)
	cutoff (from ctx0-3 memory)

Infinite Light

LLU Calculations:

$$\begin{aligned} \text{Ca} &= 1.0 \\ \text{Cd} &= n^*1 \\ \text{Cs} &= (n^*h)^{\text{power}0} \end{aligned}$$

35 Local Light

LLU Calculations:

40 att = 1.0/(K*D)
 Ca = att
 Cd = att*(n!l)
 Cs = att*((n*h)^power0)

Spot Light

45

LLU Calculations:

50 att = (s^power1)/(K*D)
 Ca = att
 Cd = att*(n^l)
 Cs = att*((n*h)^power0)

As set forth above, the mode bits controlling the vertex sequencer might not necessarily be changed by the vertex data itself or by any results derived from vertex data. To allow vertex data to modify vertex processing, LLU 1622 employs a flag register 1623 is provided. Setting bits to TRUE in this flag register allows clamping to 0.0 of calculation results if a flag is specified in the output control of the calculation. Another use of the flag register 1623 would be in setting a write mask for register writes.

The flag register 1623 is provided in LLU 1622 for performing the if/then/else clamping to 0.0 in the lighting equations at no performance penalty. The sign bit of various operands might set the flags. Table 16B illustrates the manner in which the flags in flag register 1623 are set and the resulting clamping.

Table 16B

Infinite Light

LLU Calculations:

Dflag = sign bit of (n^*l)
Sflag = sign bit of (n^*h)

Clamp:

$Ca = (0 \quad) ? 0 : Ca;$
 $Cd = (Dflag \quad) ? 0 : Cd;$
 $Cs = (Dflag | Sflag) ? 0 : Cs;$

Local Light

LLU Calculations:

Rflag = sign bit of (range-d)
Dflag = sign bit of (n^*l)
Sflag = sign bit of (n^*h)

Clamp:

$Ca = (Rflag \quad) ? 0 : Ca;$

Cd = (Rflag | Dflag) ? 0 : Cd;
Cs = (Rflag | Dflag | Sflag) ? 0 : Cs;

Spot Light

5

LLU Calculations:

Cflag = sign bit of (s-cutoff)
Rflag = sign bit of (range-d)
Dflag = sign bit of (n*l)
Sflag = sign bit of (n*h)

Clamp:

15 Ca = (Cflag | Rflag) ? 0 : Ca;
Cd = (Cflag | Rflag | Dflag) ? 0 : Cd;
Cs = (Cflag | Rflag | Dflag | Sflag) ? 0 : Cs;

Figure 21 is an illustration of the organization of the flag register 1623

20 associated with lighting module 54 of Figure 16 in accordance with one embodiment
of the present invention. The flag register 1623 contains 8 one bit flags and are set
by the sign bit of the ALU (IFLAG) or MAC0 (MFLAG) outputs.

25 When LLU 1622 outputs a scalar value to MLU 1610 where it gets smeared
into a tri-word, it specifies a mask for the flag register. If the register & mask is true,
0.0 replaces the output. Table 17 illustrates the various flags of Figure 21 to be used
in outputting ambient, diffuse, and specular attributes.

Table 17

30

Ambient Mask: C,R, U
Diffuse Mask: D, C,R, U
Specular Mask: D,S,C,R,T,U

35 The approximation used for the specular term can go negative where the
actual cos (theta)**n would go to 0.0. As a result, it is necessary to perform a
clamping operation. For this, the T, U flags are used. Table 18 illustrates various
operations of which a functional logic unit(FLU) 1621 of LLU 1622 is capable.

Note Figure 20.

40

Table 18

ZFLU_INV	$o = 1/a$	(mantissa accuracy - 12 bits)
ZFLU_ISQ	$o = 1/\sqrt{a}$	(mantissa accuracy - 6 bits)
ZFLU_PASS	$o = a$	
ZFLU_PASS1	$o = 1.0$	
ZFLU_MIN1	$o = (a < 1.0) ? a : 1.0$	
ZFLU_NOP	$o = 0.0$	

5

10

Figure 22 is an illustration of the micro-code fields associated with lighting module 54 of Figure 16 in accordance with one embodiment of the present invention. As shown, the micro-code of lighting module 54 is arranged into 33 fields making up a total width of 85 bits. Fields are delayed to match the data flow of the units. The MLU operations are done at a delay of zero, ALU operations are done at a delay of one, and RLU, LLU output operations are done at a delay of two. Each delay is equivalent to three cycles.

20 Figure 23 is a schematic diagram of sequencer 2300 associated with lighting module 54 of Figure 16 in accordance with one embodiment of the present invention. As shown, sequencer 2300 of lighting module 54 includes an input buffer 2302 adapted for receiving mode bits 202 which are indicative of the status of a plurality of modes of process operations. Also included is memory 1502 capable of storing code segments that each are adapted to carry out the process operations in accordance with the status of the modes.

25 A sequencing module 2306 is coupled between memory 1502 and buffer 2302 for identifying a plurality of addresses in memory 1502 based on a control vector 2305 derived from the mode bits. The sequencing module 2306 is further adapted for accessing the addresses in memory 1502 for retrieving the code segments that might be used to operate lighting module 54.

The sequencer **2300** of lighting module **54** is similar to that of transform module **52**. In operation, sequencer **2300** of lighting module **54** steps through a threaded control vector that is derived from threaded mode bits **202** and executes every ROM address whose corresponding control vector bit is set to “1”. The
5 control vector has the same number of bits as the ROM has words. The sequencer **2300** can step through an arbitrary control vector at the rate of a single “1” or enabled bit per a predetermined number of cycles for every thread. Commands that do not use mode bits **202** are executed by on-the-fly micro-code generation. The main difference between sequencer **2300** of lighting module **54** and sequencer **1200**
10 of transform module **52** is that sequencer **2300** of lighting module **54** can loop back and execute the lighting code up to eight times.

The sequencer **2300** of lighting module **54** has a light counter that starts at zero for each new vertex and increments by one at the end of the micro-code sequence. If the LIS field of mode bits **202** contains a “1” in the matching bit field,
15 sequencer **2300** goes back and starts over at the beginning of the lighting micro-code block. This continues until a zero is found in the LIS field or eight lights have been done. Color accumulation is done by incrementing (per light) the ALU registers that store the diffuse and specular color. Automatic memory address indexing is done
20 using the light counter to fetch the correct parameters for each light.

Figure **24** is a flowchart delineating the method by which the sequencers of the transform and lighting modules **52** and **54** are capable of controlling the input and output of the associated buffers in accordance with one embodiment of the
25 present invention. As shown, vertex data is initially received in a buffer of a first set of buffers in operation **2420**. The buffer in which the vertex data is received is based on a round robin sequence.

Subsequently, in operation **2422**, an empty buffer of a second set of buffers is
30 identified also based on a round robin sequence. The transform module **52** is

coupled between the first set of buffers and the second set of buffers. When the empty buffer of the second set of buffers is identified, the vertex data is processed in transform module and outputted from transform module to the identified empty buffer of the second set of buffers. Note operations **2424** and **2426**.

5

Similarly, an empty buffer of a third set of buffers, or slots or spaces in memory, are identified based on a round robin sequence in operation **2428**. The lighting module **54** is coupled between the second set of buffers and the third set of buffers. When the empty buffer of the third set of buffers is identified, the vertex data is processed in the lighting module, as indicated in operation **2430**. The vertex data is subsequently outputted from lighting module **52** to the identified empty buffer of the third set of buffers. See operation **2432**. It should be noted that the number of buffers, or slots in memory, is flexible and might be changed.

10

Figure **25** is a diagram illustrating the method by which the sequencers of the transform and lighting modules **52** and **54** are capable of controlling the input and output of the associated buffers in accordance with the method of Figure **24**. As shown, the first set of buffers, or input buffers **400**, feed transform module **52** which in turn feed the second set of buffers, or intermediate buffers **404, 406**. The second set of buffers **404, 406** feed lighting module **54** that drains to memory **2550**.

15

In order carry out the method set forth in Figure **25**, the slots of memory **2550** and the buffers of the first and second set are each assigned a unique identifier upon initially receiving vertex data. Further, a current state of each buffer is tracked. Such state might include an allocated state, a valid state, an active state, or a done state.

20

The allocated state indicates that a buffer/slot is already allocated to receive an output of the previous graphics-processing module, i.e. transform module or lighting module. When a write pointer is scanning the buffers/slots in the round

25

robin sequence, a buffer/slot in the allocated state cause such write pointer to increment to the next buffer or slot.

If a buffer/slot is in the valid state, the buffer/slot is available for receiving
5 vertex data. On the other hand, the active state indicates that a buffer/slot is currently in an execution state, or receiving vertex data. This active status is maintained until a thread is done after which a read pointer increments, thus placing the buffer/slot back in the valid state. It should be noted that the first set of buffers
400 are only capable of being in the valid state since there is no previous graphics-
10 processing module to allocate them.

An example of a sequence of states will now be set forth. Upon receiving vertex data in one of the first set of buffers 400 and a new set of command bits 200, such buffer is placed in the valid state, after which one of the second set of buffers
15 402, 404 is placed in the allocated state in anticipation of the output of transform module 52.

If none of the second set of buffers 404, 406 is available for allocation, the vertex data in the buffer of the first set 400 can not be processed. Further, a check
20 might be done to determine whether the code segments to be executed will interfere with any other code segments that are to be simultaneously run. If so, the vertex data in the buffer of the first set 400 will not be processed and a stall condition initiated.

25 After one of the second set of buffers 404, 406 is placed in the allocated state, the buffer of the first set 400 is placed in the active state. When transform module 52 is finished execution, the buffer of the second set 404, 406 is read and then placed in the valid state. These state changes are similarly executed during the transfer of vertex data between the second set 404, 406 and the slots of memory
30 2550.

Figure 25B illustrates the rasterizer module 56 that comprises a set-up module 57 and a traversal module 58. The rasterizer module 56 is adapted for performing area-based rasterization in an alternating manner. In particular, a plurality of polygon-defining sense points are positioned on or near the primitive after which line equations are evaluated at the points to determine which pixels reside in the primitive. During operation, this evaluation is repeated as the points are moved in an alternating manner for efficiency purposes. Further, the rasterizer module 56 might be adapted to operate without any clipping procedure.

Figure 26 illustrates a schematic of the set-up module 57 of rasterization module 56. As shown, the set-up module 57 includes a control section 61 that handles routing data and control signals to their appropriate functional units in order to perform the desired floating-point calculations. The primitive sequencer 62 handles turning sequences of vertices into triangles, lines or points. Further, floating point data path section 64 includes the multiplexers and floating point computation units that perform the math required in the set-up unit.

With continuing reference to Figure 26, output formatting section 63 handles converting the internal floating point format of edge slopes and edge values into integer formats suitable for the rasterizer since the rasterizer operates only with integer values. Of course, in alternate embodiments, the rasterizer might use a floating point thus obviating the need for output formatting section 63.

In operation, output formatting section 63 executes a block floating point conversion. As is well known, with a given number, i.e. 2.34×10^10 , floating point format tracks a mantissa (2.34) and an exponent (10) thereof. Block floating point conversion essentially manipulates the decimal place of the mantissas of incoming data such that the exponents are the same. To this end, the exponent need not be handled in rasterizer module 56.

Figure 26A is an illustration showing the various parameters calculated by set-up module 57 of rasterizer module 56 of Figure 25B. Such parameters are required for rasterizer module 56 to perform the associated functions. Upon receipt of a primitive 2600, set-up module 57 calculates three values including slopes 2601 of the primitive 5 2600, a starting position 2602 and a starting value 2604.

The slopes 2601 are used to generate coefficients for line equations of the edges of the primitive 2600 to be used during rasterization. The slopes 2601 might, for example, be calculated by using equations #4 and #5 shown below.

10

Equations #4 and #5

$$\text{slope}_A = y_0 - y_1$$

$$\text{slope}_B = x_1 - x_0$$

15

where y_0, y_1 and x_0, x_1 are coordinates of vertices shown in Figure 26A.

It should be noted that the slopes might also be calculated using the coordinates of the vertices by employing a simple rotation operation or the like.

20

The starting position 2602 indicates a starting point for area rasterization that will be set forth hereinafter in greater detail. The starting value 2604 is equal to the area of the shaded triangle shown in Figure 26A and is also used during the area-based rasterization process. Such starting value 2604 is selected so that stepping the raster 25 position about the screen while adding the slope at each step will equal zero exactly when the raster position is on the edge. Calculation of the starting value 2604 might be accomplished using Equation #6 below:

Equation #6

30

$$\text{starting_value} = \text{slope}_A * (x_s - x_0) + \text{slope}_B * (y_s - y_0)$$

where x_s, y_s = starting position **2602**

$slope_A, slope_B$ = slopes of one of the edges based on coordinates of vertices shown in Figure **26A**

5 x_0, y_0 = coordinates of one of the vertices of the edges shown in Figure **26A**

It should be understood that the foregoing values might also be calculated for other types of primitives. For example, in the case of a line, an extra slope must
10 be calculated for the four-sided bounding box. Such slope can be easily calculated by taking the reciprocal of the slope of an opposite side of the bounding box. In addition to the extra slope calculation, it is noted that another starting value needs to be calculated in the case of the line primitive.

15 Figure **27** illustrates the method by which rasterizer module **56** handles one of a plurality of primitives, e.g. triangles. In particular, an initial operation is first performed by set-up module **57** of rasterizer module **56**. Upon receipt of a primitive,
line equation coefficients of line equations are determined for lines that define the primitive in operation **2700** using slopes **2601** of Figure **26A** in a manner that is well
20 known to those with ordinary skill in the art. As is well known, three line equations are required to define a triangle. On the other hand, a primitive such as a line is drawn as a rectangle or parallelogram with four sides and four line equations.

Thereafter, in operation **2702**, the line equation coefficients are modified if
25 any primitive vertex(es) has a negative W-coordinate. Additional information regarding this process will be set forth hereinafter in greater detail with reference to Figure **32**.

It should be noted that set-up module **57** of rasterizer module **56** also
30 computes a bounding box of the primitive. For most triangles, the bounding box includes the minimum and maximum values of the three vertextes. For lines, the

four parallelogram corners of the bounding box are calculated. For triangles or lines that have a vertex with a negative W-coordinate, an area that is to be drawn extends beyond the convex hull of the vertices.

5 One of the commands of OpenGL® is a scissor rectangle which defines a boundary outside of which is not to be drawn. The set-up module 57 of rasterizer module 56 calculates the intersection of the bounding box and the scissor rectangle. Since the scissor rectangle is a rectangle, four additional line equations are afforded. It should be noted that the line equations associated with the scissor rectangle have a
10 trivial form, i.e. horizontal or vertical.

Furthermore, in 3-D space, the near plane and far plane are parallel and at right angles to the line of sight. In the case of the primitive being a triangle, three vertexes are included which define a plane that might have any orientation. The
15 intersections of the plane of the primitive and the near and far planes include two lines with two associated line equations.

Accordingly, each primitive has a total of nine or ten line equations depending on whether it takes the form of a triangle or a line, respectively. Again, in
20 the case of the triangle, such line equations include the three line equations which define the triangle, the four line equations defining the bounding box and the two line equations which define the intersections of the plane in which the primitive resides, and near and far planes.

25 With continuing reference to Figure 27, the process progresses in operation 2704 by positioning a plurality of points on or near the primitive. The starting position 2602 dictates such positioning, as shown in Figure 26A. Such points define an enclosed convex region and reside at corners of the convex region. Figure 27A illustrates such sense points 2705 that enclose convex region 2707, e.g. a rectangle.
30 In one embodiment, such rectangle might be 8x2 pixels in size. Further, the points

might be initially positioned to enclose a top vertex of the primitive. As an option, this might be accomplished using truncation.

Once the primitive is positioned, the process is continued by traversal module **58** which begins in operation **2706** by processing rows of the primitive in a manner set forth below. After the processing of each row, it is determined whether a jump position has been found in decision **2708**. A jump position is a starting position for processing the next row and will be described hereinafter in greater detail. If it is determined in decision **2708** that a jump position has been found, the sense points that define the convex region are moved thereto in operation **2710**. If, however, it is determined that a jump position has not been found, the process is ended. It should be noted that, in an alternate embodiment, columns, diagonals or any other type of string might be processed in operation **2706** instead of rows.

Figure **28** is a flowchart illustrating a process of the present invention associated with the process row operation **2706** of Figure **27**. As shown, the process begins by computing the sense points in operation **2800** in order to determine whether the polygon-defining sense points might be moved right in decision **2801**. Such decision is made based on the position of the rightmost sense points. If the rightmost sense points are not positioned outside the same edge or edges of the primitive, rightward movement is permitted and a position (X and Y coordinates) to the right of the current position is stored as a snap location in operation **2802**. If, however, both rightmost sense points are positioned outside one or more edges of the primitive, rightward movement is not permitted and operation **2802** is skipped.

Next, the line equations are evaluated at the points of the convex region, e.g. rectangle, in operation **2804**. The evaluation includes determining if the points reside in the primitive. Such determination as to whether the points reside in the primitive might include determining whether the evaluation of each of the line equations renders a positive value or a negative value at each of the sense points.

The line equations can be formulated to be positive inside the primitive and negative outside. Inclusive edges, for which pixels that lie exactly on the edge should be drawn, evaluate to zero and might be treated as positive. Exclusive edges, which should not be drawn, can be made negative by initially subtracting a value of 5 one from the starting line equation value. Thus pixels on exclusive edges evaluate to a negative value (-1) instead of a positive zero. This permits the sense point interpretation to ignore the inclusive/exclusive policy and just test the line equation sign.

10 After the line equations are evaluated at the points, it is determined whether a current position of the sense points constitutes a jump position in decision **2806**. It should be noted that a jump position is stored only if the two bottom sense points are not both outside an edge. If it is determined in decision **2806** that a jump position has been found, the jump position is calculated and stored (or replaces a previously 15 stored jump position if existent) in operation **2808**. If not, however, operation **2808** is skipped.

With continuing reference to Figure 28, it is then determined in decision 20 **2810** whether leftmost sense points are both outside an edge of the primitive. Again, this process entails determining whether the evaluation of the line equations at both of the leftmost sense points renders positive or negative values. In particular, upon computation of the coefficients of the nine or ten edge equations at the pertinent sense points, nine or ten values are rendered that have nine or ten sign bits. To determine if the current side is completely outside any edge, for example, the present 25 invention AND's the ten sign bits from the two sense points together. If any bit(s) survive, then both points are outside that edge.

If it is determined that the leftmost sense points are not both outside an edge of the primitive, it is concluded that there still remains further portions of the 30 primitive to be considered in the leftward direction, and the sense points are moved left in operation **2812**. If it is determined in decision **2810** that both leftmost sense

points are indeed outside the edge of the primitive, it is concluded that there no longer remains further portions of the primitive to be considered in the leftward direction. Next, in decision **2814**, it is determined whether there is a snap location that resulted from operation **2802**.

5

If it is determined in decision **2814** that a snap location does not exist, the process is done. If, however, a snap location does exist, the sense points are moved to the snap location in operation **2816**. Thereafter, operations similar to those of operations **2804-2812** are executed to map a right side of the primitive. This begins 10 in operation **2818** by the line equations being evaluated at the points of the convex region.

After the line equations are evaluated at the points, it is determined whether a current position of the sense points constitutes a jump position in decision **2820**. If 15 it is determined in decision **2806** that a jump position has been found, the jump position is calculated and stored in operation **2822**. If not, however, operation **2822** is skipped.

With continuing reference to Figure **28**, it is then determined in decision 20 **2824** whether rightmost sense points are both outside an edge of the primitive. If it is determined that the rightmost sense points are not both outside an edge of the primitive, it is concluded that there still remains further portions of the primitive in the rightward direction to be considered, and the sense points are moved right in operation **2826**. If it is determined in decision **2824** that both rightmost sense points 25 are outside the edge of the primitive, it is concluded that there no longer remains further portions of the primitive to be considered in the rightward direction, and the instant process is done.

Figures **28A** and **28B** are illustrations of the sequence in which the sense points 30 of the present invention might be moved about the primitive **2850**. It should be noted

that various alterations might include determining whether the points can go left in decision 2800 and proceeding right initially. Further, the line equations might be defined to indicate whether the points are inside or outside the primitive in any desired way.

5

To avoid stepping in a repeating loop, the present invention thus employs an overall direction of movement during rasterization. The initial implementation proceeds top-down, visiting every convex region on a row before stepping down to the next. By processing rows top-down as well as never stepping right then left or left then right, 10 loops are thus avoided.

An example of the foregoing process might be shown with reference to the polygon-defining points, P1, P2, P3 and P4 of Figure 27A. In operation, pairs of adjacent sense points can be examined to determine whether stepping in their direction 15 would be productive. For example, if both P3 and P4 in Figure 27A were outside of an edge of a polygon, but P1 and/or P2 are not, then clearly the drawable inside region lies to the left, not to the right. Thus the sense points should not move right. Conversely, if both P3 and P4 are inside all the edges, then there is a drawable area just beyond P3 and P4, and stepping right is appropriate. Indeed, if P3 and P4 were not outside the same 20 edge or edges, stepping right would be productive. This same logic applies to stepping upwards guided by P1 and P3, or stepping left guided by P1 and P2, or stepping downwards based on P2 and P4.

The foregoing process thus moves, or steps, the convex region defined by the 25 points around the inside of the primitive, using sense points as a guide. Since the convex region defined by the points might be large, many pixels might be tested simultaneously. During use, if all sense points are inside all edges of the primitive, then all the enclosed pixels must be drawable (assuming a convex primitive). A significant advantage is afforded by testing the corners, namely the ability of proving an arbitrary 30 area of the primitive is inside, outside or split. Only in the latter case do the individual pixels in the convex region defined by the points need to be tested. In such case, the

pixels in the convex region defined by the points might be tested one-by-one or by another method in order to determine whether they reside in the primitive. Furthermore, the sense points might reduce the amount of further testing required by defining which edges(s) split the area and which do not.

5

Figure 29 is a flowchart illustrating an alternate boustrophedonic process of the present invention associated with the process row operation 2706 of Figure 27. As shown, it is first determined in decision 2900 whether a previous movement was in a first or second direction. If there was not any actual previous movement, a default 10 previous movement might be assumed. If it is determined in decision 2900 that the previous movement was in a second direction, the line equations are evaluated at the points of the convex region, e.g. a rectangle, in operation 2902 in a manner similar to operation 2804 of Figure 28.

15 With continuing reference to Figure 29, it is subsequently determined in decision 2904 as to whether sense points of a first side of the rectangle are both outside an edge of the primitive. If not, the sense points are moved or stepped in the first direction in operation 2906. Upon it being determined that the sense points of the first side of the rectangle are both outside an edge of the primitive, it is then determined in decision 20 2905 whether the points can be moved downwardly or, in other words, whether the current position constitutes a jump position. If so, a jump position is calculated and stored in operation 2908 after which the process is done.

25 On the other hand, if it is determined in decision 2900 that the previous movement was in a first direction, operations similar to those of operation 2902-2908 are carried out. In particular, the line equations are evaluated at the points of the convex region, e.g. a rectangle, in operation 2910. It is then determined in decision 2912 as to whether sense points of a second side of the rectangle are both outside an edge of the primitive. If not, the sense points are moved or stepped in the second direction in 30 operation 2914. Upon it being determined that the sense points of the second side of the

rectangle are both outside an edge of the primitive, it is then determined in decision 2913 whether the points can be moved downwardly or, in other words, whether the current position constitutes a jump position. If so, a jump position is calculated and stored in operation 2916 after which the process is done.

5

Figure 29A is an illustration of the sequence in which the sense points of the present invention are moved about the primitive in accordance with the boustrophedonic process of Figure 29. The foregoing boustrophedonic rasterization constrains the sequence to obey certain rules that offer better performance for hardware. As shown, 10 the boustrophedonic rasterization affords a serpentine pattern that folds back and forth. A horizontal boustrophedonic sequence, for example, might generate all the pixels within a primitive triangle that are on one row from left to right, and then generate the next row right to left, and so on. Such a folded path ensures that an average distance from a generated pixel to recently previously generated pixels is relatively small.

15

Generating pixels that are near recently previously generated pixels is important when recent groups of pixels and/or their corresponding texture values are kept in memories of a limited size. The boustrophedonic sequence more often finds the pixels or texture values already loaded into such memories, and therefore 20 repeating the memory load occurs less often.

As an option, at least one boundary might be used which divides the primitive into a plurality of portions prior to rasterization. In operation, the points might be moved in each of the portions separately. Further, the points might be 25 moved through an entirety of a first one of the portions before being moved in a second one of the portions.

Figure 30 is a flowchart illustrating an alternate boustrophedonic process using boundaries. As an option, the decision whether to use boundaries might be 30 based on a size of the primitive. As shown in Figure 30, the boustrophedonic process which handles boundaries is similar to that of Figure 27 with the exception

of an additional operation **3000** wherein at least one boundary is defined which divides the primitive into a plurality of portions or swaths.

With continuing reference to Figure 30, an additional decision **3001** follows
5 the completion of every portion of the primitive. In particular, it is determined in decision **3001** whether a start position of an adjacent portion was found in operation **3006**. If so, the convex region defined by the sense points is moved to a start position of an adjacent portion of the primitive in operation **3002** and operations **3004-3010** are repeated for the new portion of the primitive. Further information
10 relating to the determination of the start position in operation **3006** will be set forth in greater detail during reference to Figure 31.

Figure **31A** is an illustration of the process by which the convex region of the present invention is moved about the primitive in accordance with the boundary-based boustrophedonic process of Figure 30. As shown, the first portion that is processed is that which includes the topmost vertex of the primitive. During operation, a left neighboring portion is processed after which the adjacent left neighboring portion is processed and so on. This is continued until there are no remaining left neighboring portions. Next, a neighboring portion to the right of the
15 first portion is processed after which the adjacent right neighboring portion is processed and so on until all of the right neighboring portions are processed. It should be appreciated that other types of ordering schemes might be utilized per the desires of the user.
20

Figure **31** is a flowchart showing the process associated with the process row operation **3006** of Figure 30. Such process is similar to the boustrophedonic process of Figure 29 with the exception of decisions **3118** through **3121**. Decisions **3118** and **3120** both determine whether any of the sense points have passed any boundary. Only if it is determined that the sense points are still within the boundaries is the
25 respective loop continued.
30

In operations **3119** and **3121**, starting positions of adjacent portions of the primitive are sought and stored when it is determined in decisions **3118** and **3120** that any sense points of the convex region have passed any boundary, respectively.

5 As shown in Figure **31A**, such starting positions **3126** are each defined as being the topmost point of a portion of the primitive existent beyond a boundary. By storing this position, a starting point is provided when the process is repeated for the adjacent boundary-defined portion of the primitive.

10 It should be noted that operations **3119** and **3121** are both performed while processing the first portion of the primitive. While not expressly shown in Figure **31**, only a first one of such operations is performed when processing portions to the left of the first portion, while only a second one of such operation is performed when processing portions to the right of the first portion. In other words, when processing portions to the left of the first portion, starting positions are only determined when a leftmost boundary of the currently processed portion has been exceeded. Similarly, when processing portions to the right of the first portion, starting positions are only determined when a rightmost boundary of the currently processed portion has been exceeded.

15

20 Using boundaries during rasterization solves a very critical problem during pipeline processing. If a primitive is very wide, the storage associated with the pixels of a single row might not fit in a limited-size memory. Rasterization with boundaries divides the triangle into limited-width rows (or columns), and generates all the pixels within such a portion before moving on to the next portion.

25

For example, even if a triangle is 100 pixels wide, a limited-size pixel or texture memory might only hold information for the previous 20 pixels.

Constraining the pixel sequence to stay within ten-pixel-wide vertical portions
30 allows all the pixels on the previous and current rows to fit in the memory. This means that a boustrophedonic sequence within a boundary-defined portion would

always have the previous pixel on the current row (if any) in the memory, as well as the pixels in the row above (if any) in the memory as well.

Most underlying memory systems transfer blocks of data with a certain overhead per block. Small accesses to the memory system are penalized heavily by this overhead. In order to be efficient, larger accesses are employed and the rest of the block is maintained in case it might be used next. Beyond that, a cache memory system keeps a plurality of these recent blocks, increasing the probability that memory accesses can be avoided.

10

The boustrophedonic sequence of the present invention exploits the single-retained-block concept when it reverses and handles pixels immediately below one end of the current line. Further, the boustrophedonic sequence exploits cache when it limits rasterization to portions of a particular size. Specifically, two scanlines within a portion should fit in the cache, so throughout the second scanline, benefits might be incurred from cache storage of the first scanline.

15

There is no constraint on the sequence or number of boundary-defined portions. Although the present description uses the example of vertical portions and a horizontal boustrophedonic pattern, similar principles might extend to horizontal portions, vertical boustrophedonic patterns or even to diagonal portions and patterns. In one embodiment, the length of the strings (e.g. rows, columns, diagonals, etc.) might be each limited to be less than a dimension of the primitive along which the string resides.

20

Figure 32 is a flowchart showing the process associated with operation 2702 of Figure 27. The instant process is designed to handle a primitive with portions that reside behind the eye. These outlying portions might cause problems in subsequent rasterization operations. To accomplish this, the instant process employs a variable, W that is commonly used for projection i.e., for viewing objects in perspective. The variable W is a number that the other coordinates, X, Y and Z, are divided by in

order to make nearby things larger and far things smaller. The variable W is representative of a distance between a center of projection and the corresponding vertex.

5 As shown in Figure 32, a primitive is first received that is defined by a plurality of vertices. Each of such vertices includes a W-value. Upon the receipt of the primitive, the set-up module serves to define lines that characterize the primitive based on the vertices. Note operation 3200.

10 The W-values are then analyzed in decision 3202. As shown, if one of the W-values is negative, a line equation for a line opposite the vertex having the negative value is flipped in operation 3204. In other words, the coefficients of the line equation are multiplied by -1. Further, if two of the W-values are negative, line equations for lines connecting the vertex having a positive W-value and each of the vertexes having
15 negative W-values are flipped in operation 3206. If three of the W-values are negative, a cull condition 3207 occurs where the present invention culls the triangle. Still yet, if none of the W-values are negative, no additional action is taken.

20 Figures 32A - 32C illustrate the manner in which flipping line equations affects which portion of the screen is processed. Figure 32A shows the case where none of the W-values are negative and the line equations are left unaltered. As shown, an interior portion of the primitive is filled in such case.

25 Figure 32B shows the case where one of the W-values is negative and which of the line equations is flipped accordingly. As shown, the portion of the primitive opposite the vertex is filled in the present case. In particular, the area to be drawn is bounded by two lines that are co-linear with the two triangle sides sharing the -W vertex, and further bounded by a side of the triangle that shares the two +W vertexes.

30 Figure 32C shows the case where two of the W-values are negative and which of the line equations are flipped accordingly. As shown, the portion of the primitive

opposite the vertexes is filled using the methods and/or processes set forth hereinabove with reference to Figures 27-32. In other words, the area to be drawn is bounded by two lines that are co-linear with the two triangle sides sharing the +W vertex, and further contiguous to the +W vertex.

5

The present invention is thus capable of handling all three of the foregoing cases. If part of the triangle is beyond the near and/or far plane, it draws only the portion within those planes. If the triangle has one or two negative Z vertexes, only the correct +Z portion is drawn.

P0002600-P0002600-P0002600-P0002600-P0002600

Even if all vertexes are off-screen, and the triangle extends from behind the eye to beyond the far plane, whatever pixels are inside the triangle and on the screen and have Z between the near and far limits. The present invention ensures that little time is wasted exploring bad pixels. This is possible because all clipping, by screen 10 edge or the near or far plane, always results in a convex region on-screen which can be explored easily.

A problem sometimes arises when the starting point is not inside the area to be filled. This can occur if the top vertex is off-screen or is clipped by the near or far 15 plane. In this case, the traversal stage must search for the top point of the drawn region, starting from above. It can do this efficiently by being guided by the signs of the triangle edge slopes and the Z slope. It can test the triangle line equations to discover it is outside the drawn region and why. When it knows what edge(s) and/or Z limit it is outside of, it knows what direction(s) to step that brings it closer to that 20 edge or limit. By moving horizontally in preference to vertically (when there is a choice), searching for the drawn region guarantees it finds the top drawable pixel if there is one. This problem also occurs with external (-W) triangles that open up. In this case, the drawn area extends above all three vertexes.

25 In one embodiment of the present invention, traversal proceeds from top to bottom of the triangle. The starting point is the top vertex of the triangle if none

have a negative W-value and the top vertex is in the scissor rectangle. Otherwise, a point on the top of the scissor rectangle is chosen. Since traversal always begins within the scissor rectangle and never ventures out of it, only the portion of the triangle within the scissor rectangle is ever drawn, even if the area enclosed by the edges extends far beyond the scissor rectangle. In this way, simple scissor rectangle-edge clipping is effected.

While various embodiments have been described above, it should be understood that they have been presented by way of example only, and not limitation. Thus, the breadth and scope of a preferred embodiment should not be limited by any of the above-described exemplary embodiments, but should be defined only in accordance with the following claims and their equivalents.